Deep Learning in Asset Pricing

Luyang Chen†, Markus Pelger‡, Jason Zhu‡

† Institute for Computational and Mathematical Engineering, Stanford University

‡ Department of Management Science and Engineering, Stanford University

Contact Information:
Markus Pelger
Stanford University, Huang 312
475 Via Ortega, Stanford, CA 94305

Phone: +1 (510) 859 5151
Email: mpelger@stanford.edu

Stanford University

Stanford ENGINEERING | Advanced Financial Technologies Laboratory

Abstract

We estimate a general non-linear asset pricing model and optimal portfolios with deep neural networks applied to all U.S. equity data combined with a substantial set of macroeconomic and firm-specific information. Our crucial innovation is the use of the no-arbitrage condition as part of the neural network algorithm. We estimate the stochastic discount factor (SDF or pricing kernel) that explains all asset prices from the conditional moment constraints implied by no-arbitrage. For this purpose, we combine three different deep neural network structures in a novel way: A feed-forward network to capture non-linearities, a recurrent (LSTM) network to find a small set of economic state processes, and a generative adversarial network to identify the portfolio strategies with the most unexplained pricing information. Our SDF is a portfolio of all traded assets with time-varying portfolio weights which are general functions of the observable firm-specific and macroeconomic variables. Our model allows us to understand what are the key factors that drive asset prices, identify mis-pricing of stocks and generate the optimal portfolio with the highest risk-adjusted return. Empirically, our optimal portfolio strategies have a risk-adjusted return out-of-sample (annual Sharpe-ratio 2.1) that outperforms all other benchmark models in the literature.

Introduction

Machine Learning in Finance

In efficient markets, asset returns are dominated by unforecastable news. Hence, financial returns exhibit a very low signal-to-noise ratio which distinguishes them from other applications of machine learning. In this paper we include financial constraints (no-arbitrage) in the learning algorithm to significantly improve the signal and obtain considerably better results than with off-the-shelf machine learning algorithms.

Research questions:

1. Explain asset prices for different assets
2. Design optimal risk-adjusted portfolios
3. Find mis-priced assets to earn alpha
4. Use all available information in the market and understand which information is relevant

This Paper

We estimate the pricing model with deep neural networks. The crucial innovation is to include the no-arbitrage condition in the neural network algorithm and combine three neural network structures in a novel way:

1. Non-linearity: Feed-forward network captures non-linearities
2. Time-variation: Recurrent (LSTM) network finds a small set of economic state processes
3. Pricing all assets: Generative adversarial network identifies the states and portfolios with most unexplained pricing information
4. Dimension reduction: Regularization through no-arbitrage condition
5. Signal-to-noise ratio: No-arbitrage conditions increase the signal-to-noise ratio

⇒ General model that includes all existing models as a special case

Contribution

1. Empirically outperforms all benchmark models in terms of out-of-sample Sharpe, explained variation and pricing errors.
2. Optimal portfolio has out-of-sample annual Sharpe-ratio of 2.1.
3. Non-linearities and interaction between firm information matters.
4. Most relevant firm characteristics are price trends, liquidity, and value.

No-Arbitrage Asset Pricing

The Stochastic Discount Factor (SDF)

The key object in No-Arbitrage Pricing Theory (APT) is the stochastic discount factor (SDF, also called pricing kernel or equivalent martingale measure) which explains differences in risk and asset prices. It is defined as a stochastic process $\{M_t\}$ such that for any security with payoff x_{t+1} at time $t+1$, the price of that security at time t is

$$P_t = \mathbb{E}_t[M_{t+1}x_{t+1}]$$

Equivalently the fundamental no-arbitrage condition can be expressed as

$$\mathbb{E}_t[M_{t+1}R^e_{i,t+1}] = 0$$

where $R^e_{i,t+1} = R_{i,t+1} - R_f$ is the excess return (return minus risk-free rate) of asset i at time $t+1$. Without loss of generality the SDF is the gross return of the SDF portfolio:

$$M_{t+1} = 1 + \sum_{i=1}^{N} w_{i,t}R^e_{i,t+1}$$

The SDF portfolio $-\sum_{i=1}^{N} w_{i,t}R^e_{i,t+1}$ is the mean-variance efficient portfolio with the highest conditional Sharpe-ratio.

Equivalent factor model representation

The no-arbitrage condition is equivalent to a beta representation

$$\mathbb{E}_t[R^e_{i,t+1}] = \frac{Cov_t(R^e_{i,t+1}, F_{t+1})}{Var_t(F_{t+1})} \cdot \mathbb{E}_t[F_{t+1}] = \beta_{i,t}\mu^F_t$$

with factor $F_t = 1 - M_t$. Without loss of generality no-arbitrage is equivalent to a one-factor model with time-varying loadings:

$$R^e_{t+1} = \beta_t F_{t+1} + \varepsilon_{t+1}$$

We estimate portfolio weights of the SDF factor and the risk-exposure β as general functions of all available information at time t.

Model

Loss Function (General Method of Moments)

We estimate the weights of the SDF portfolio by minimizing the fundamental no-arbitrage moment conditions. The pricing formula is equivalent to an infinite number of moment conditions. For any $\mathcal{F}_t$-measurable variable $\hat{I}_{i,t}$, the corresponding moment condition equals

$$\mathbb{E}[M_{t+1}R^e_{i,t+1}\hat{I}_{i,t}] = 0$$

which implies a loss function. The empirical loss function minimizes the sample mean pricing error (corrected for the different number of time-series observation for each asset):

$$L(\hat{I}) = \frac{1}{N}\sum_{i=1}^{N}\frac{|T_i|}{T}\Big(\frac{1}{|T_i|}\sum_{t\in T_i} M_{t+1}R^e_{i,t+1}\hat{I}_{i,t}\Big)^2$$

Model Structure

We use a "recurrent" network followed by a "feed-forward" networks to estimate the portfolio weights of the SDF factor. We refer to this network as the SDF Network. We construct conditioning variables via a Conditional Network with a similar neural architecture. The Conditional Network serves as an adversary and competes with the SDF Network to identify the assets and portfolio strategies that are the hardest to explain.

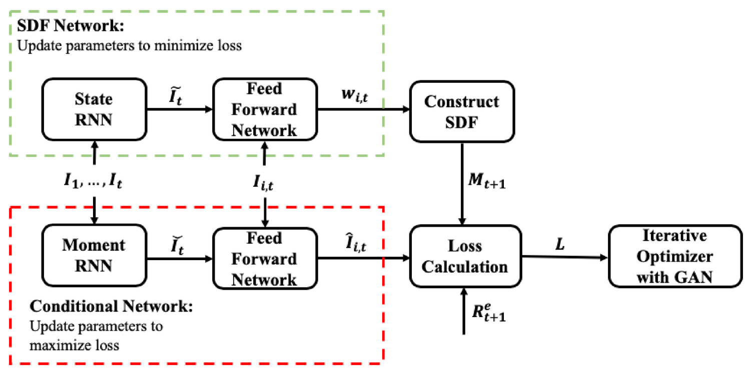

GAN - Generative Adversarial Network

SDF Network and Conditional Network play a zero-sum game. We iteratively update the two networks. For a given conditional variable $\hat{I}_t$, the SDF Network minimizes the loss. For a given SDF, the Conditional Network finds $\hat{I}_t$ with the largest loss (most mispricing). The intuition is to find the economic states and assets with the most pricing information.

Recurrent Neural Network

We use recurrent neural network (RNN) with Long-Short-Term-Memory (LSTM) cells to transform all macroeconomic time-series into a low dimensional vector of stationary state variables. This is because:

1. Economic time-series data is often non-stationary. Therefore, necessary transformation is required.
2. Asset prices depend on economic states. Therefore, simple differencing of non-stationary data is not sufficient.

Four macro state variables (LSTM outputs) are enough to capture macroeconomic movements.

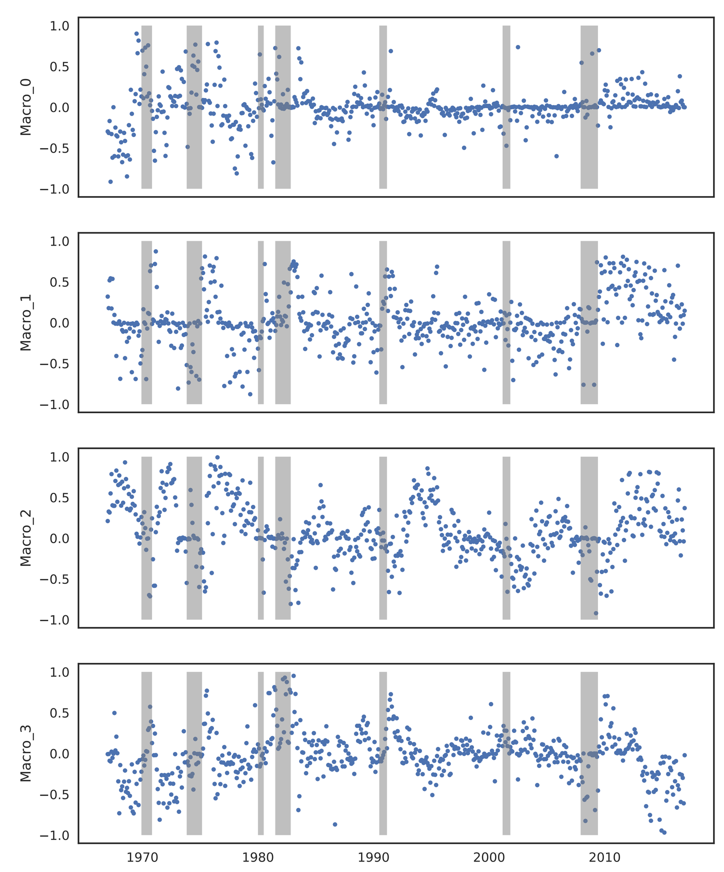

Model Evaluation

Evaluation Metrics

We evaluate the performance of our model by calculating the Sharpe-ratio of the SDF portfolio, the amount of unexplained variation and the pricing errors of the model. We obtain the systematic and non-systematic return component by projecting returns on the risk exposure β, which is estimated by fitting a feed-forward network to predict $\mathbb{E}_t[R^e_{t+1}F_{t+1}]$.

$$\varepsilon_{t+1} = \left(I_N - \beta_t(\beta_t^\top\beta_t)^{-1}\beta_t^\top\right)R^e_{t+1}$$

We calculate the following three performance metrics:

- Sharpe ratio of SDF factor (SR)

$$\frac{\mathbb{E}[F]}{\sqrt{Var(F)}}$$

- Unexplained variation (UV)

$$\frac{1}{T}\sum_{t=1}^{T}\Big[\frac{1}{|N_t|}\sum_{i\in N_t}(\varepsilon_{i,t+1})^2\Big] \Big/ \frac{1}{T}\sum_{t=1}^{T}\Big[\frac{1}{|N_t|}\sum_{i\in N_t}(R^e_{i,t+1})^2\Big]$$

- Time-weighted average pricing error (Alpha)

$$\frac{1}{N}\sum_{i=1}^{N}\frac{|T_i|}{T}\Big[\frac{1}{|T_i|}\sum_{t\in T_i}\varepsilon_{i,t+1}\Big]^2$$

Benchmark Models

We compare our GAN model with three benchmark models, including (1) Fama-French factor investing approach (FF-3 and FF-5); (2) linear asset pricing models using Least Squares (LS) and Lasso Regression (Lasso); and (3) return forecasting with feed-forward network (FFN) as described in Gu et.al. (2018).

Data

Returns and Firm Specific Characteristic Variables

We obtained monthly equity returns data for all securities on CRSP from 1967 to 2016. we constructed 46 firm-specific characteristics either defined on the Kenneth French Data Library or compiled by Freyberger et.al. (2019). All these variables are constructed from either accounting variables from CRSP/Compustat database or past returns from CRSP.

Macroeconomic Variables

We constructed a collection of 178 macroeconomic variables, which come from three sources. We take 124 macroeconomic predictors from the FRED-MD database. We also take cross-sectional median time series of the 46 firm characteristics and an additional 8 macroeconomic predictors from Welch et.al. (2007).

Empirical Results

Performance on Individual Stocks

- Our SDF portfolio has the highest risk-adjusted payoff (Sharpe-ratio) out-of-sample and lower risk as measure by drawdown:

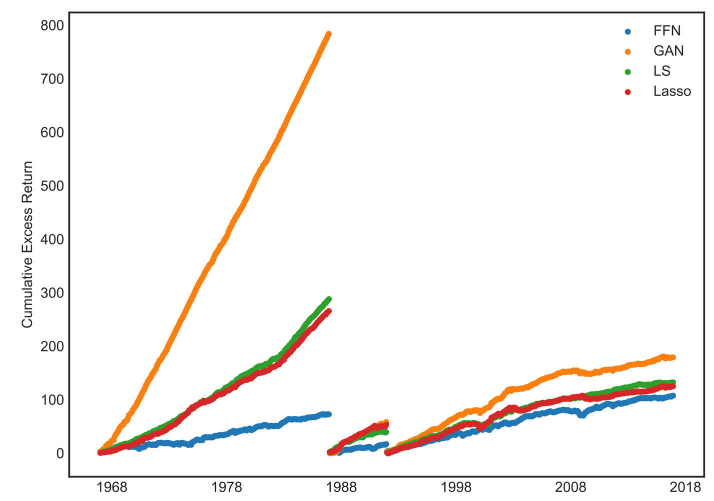

	SR (monthly)			Max Loss			Max Drawdown		
Model	Train	Valid	Test	Train	Valid	Test	Train	Valid	Test
FF-3	0.27	-0.09	0.19	-2.45	-2.85	-4.31	7	10	10
FF-5	0.48	0.40	0.22	-2.62	-2.33	-4.90	4	3	7
LS	1.20	0.65	0.44	-2.19	-2.04	-4.55	1	4	4
Lasso	1.11	0.88	0.42	-2.83	-1.85	-6.56	2	3	7
FFN	0.30	0.28	0.36	-3.88	-4.93	-4.07	7	4	5
GAN	**3.26**	**0.97**	**0.60**	**-0.09**	**-1.01**	**-4.48**	**1**	**2**	**3**

- Our model explains the variation and expected returns of individual stocks better than the other benchmark models.

	UV			Alpha ($\times 10^{-4}$)		
Model	Train	Valid	Test	Train	Valid	Test
LS	0.88	0.95	0.96	0.80	5.39	1.40
Lasso	0.86	0.94	0.95	0.80	5.30	1.33
FFN	0.84	0.93	0.94	0.84	5.46	1.35
GAN	**0.81**	**0.91**	**0.93**	**0.81**	**5.39**	**1.31**

- Conventional Fama-French factors do not span SDF factor.

	Mkt-RF	SMB	HML	RMW	CMA	intercept
coefficient	0.07	0.01	0.03	0.13	-0.01	0.38
correlation	0.14	-0.11	0.23	0.31	0.04	-

Predictive Pricing Performance

- Portfolios with higher predicted β have higher average returns. We sort stocks into decile portfolios based on their conditional β. The first decile portfolio is based on the smallest and the last on the largest decile.

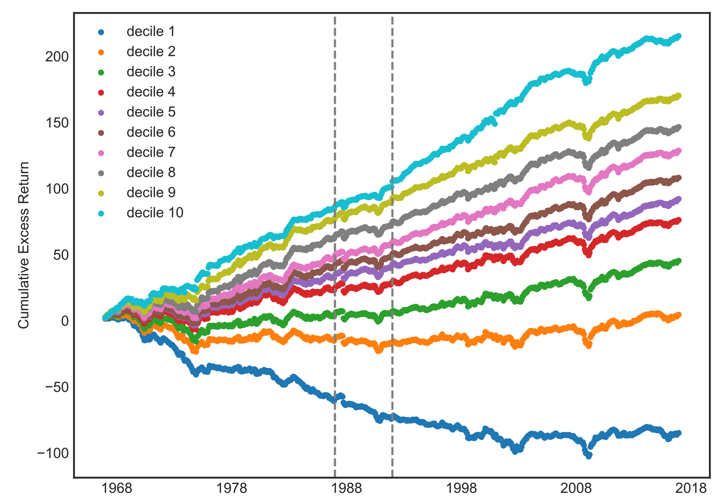

Model Interpretation and Structure of Weights

- We rank the 46 firm characteristics based on their sensitivity. The most important variables are Short-Term Reversal (ST_REV), Momentum (r12_2) and Standard Unexplained Volume (SUV).

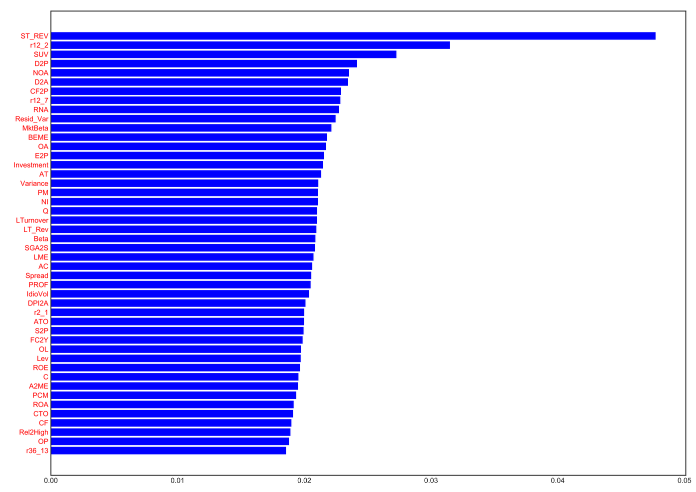

- Nonlinear interactions between characteristics are important, which cannot be captured by standard models. The plots show the non-linear relationship in the SDF portfolio weights for Short-Term Reversal (ST_REV), Momentum (r12_2) and Standard Unexplained Volume (SUV) and their interactions.